arXiv:1609.06418v1 [math.ST] 21 Sep 2016

Multiple Testing via Relative Belief Ratios

Michael Evans and Jabed Tomal
Department of Statistical Sciences
University of Toronto

Abstract: Some large scale inference problems are considered based on using the relative belief ratio as a measure of statistical evidence. This approach is applied to the multiple testing problem. A particular application of this is concerned with assessing sparsity. The approach taken to sparsity has several advantages as it is based on a measure of evidence and does not require that the prior be restricted in any way.

Key words and phrases: multiple testing, sparsity, statistical evidence, relative belief ratios, priors, checking for prior-data conflict, testing for sparsity, relative belief multiple testing algorithm.

1 Introduction

Suppose there is a statistical model $\{f_\theta : \theta \in \Theta\}$ and the parameter of interest is given by $\Psi : \Theta \to \Psi$ (we don't distinguish between the function and its range to save notation) where Ψ is an open subset of R^k. Our concern here is with assessing individual hypotheses $H_{0i} = \{\theta : \Psi_i(\theta) = \psi_{0i}\}$, namely, H_{0i} is the hypothesis that the i-th coordinate of ψ equals ψ_{0i}. If it is known that $\psi_i = \Psi_i(\theta) = \psi_{0i}$, then the effective dimension of the parameter of interest is $k-1$ which results in a simplification of the model. For example, the ψ_i may be linear regression coefficients and $\psi_{0i} = 0$ is equivalent to saying that an individual term in the regression is dropped and so simplifies the relationship.

Considering all these hypotheses separately is the multiple testing problem and the concern is to ensure that, while controlling the individual error rate, the overall error rate does not become too large especially when k is large. An error means either the acceptance of H_{0i} when it is false (a false negative) or the rejection of H_{0i} when it is true (a false positive). This problem is considered here from a Bayesian perspective. A first approach is to make an inference about the number of H_{0i} that are true (or false) and then use this to control the number of H_{0i} that are accepted (or rejected). Problems arise when k is large, however, due to an issue that arises with typical choices made for the prior. This paper considers these problems and proposes solutions.

The problem of making inferences about sparsity is clearly related. For suppose that there is a belief that many of the hypotheses H_{0i} are indeed true

but there is little prior knowledge about which of the H_{0i} are true. This is effectively the sparsity problem. In this case it is not clear how to choose a prior that reflects this belief. A common approach in the regression problem is to use a prior that, together with a particular estimation procedure, forces many of the ψ_i to take the value ψ_{0i}. A difficulty with such an approach is that it is possible that such an assignment is an artifact of the prior and the estimation procedure and not primarily data-based. For example, the use of a Laplace prior together with MAP estimation is known to accomplish this in certain problems. It would be preferable, however, to have a procedure that was not dependent on a specific form for the prior but rather worked with any prior and was based on the statistical evidence for such an assignment contained in the data. The methodology for multiple testing developed here accomplishes this goal.

The prior π on θ clearly plays a key role in our developments. Suppose that π is decomposed as $\pi(\theta) = \pi(\theta \mid \psi)\pi_\Psi(\psi)$ where $\pi(\theta \mid \psi)$ is the conditional prior on θ given $\Psi(\theta) = \psi$ and π_Ψ is the marginal prior on ψ. If π_Ψ is taken to be too concentrated about ψ_0, then there is the risk of encountering prior-data conflict whenever some of the H_{0i} are false. Prior-data conflict arises whenever the true value of the parameter lies in a region of low prior probability, such as the tails of the prior, and can lead to erroneous inferences. Consistent methods for addressing prior-data conflict are developed in Evans and Moshonov (2006), Evans and Jang (2011a) and a methodology for modifying the prior, when a conflict is detected, is presented in Evans and Jang (2011b). Generally, a check for prior-data conflict is made to ensure that the prior has been chosen reasonably, just as model checking is performed to ensure that the model is reasonable, given the observed data. The avoidance of prior-data conflict plays a key role in our developments.

The assessment of the evidence for the truth of H_{0i} is based on a measure of the evidence that H_{0i} is true as given by the relative belief ratio for ψ_i. Relative belief ratios are similar to Bayes factors and, when compared to p-values, can be considered as more appropriate measures of statistical evidence. For example, evidence can be obtained either for or against a hypothesis. Moreover, there is a clear assessment of the strength of this evidence so that when weak evidence is obtained, either for or against a hypothesis, this does not entail acceptance or rejection, respectively. Relative belief ratios and the theory of inference based on these quantities, is discussed in Evans (2015) and a general outline is given in Section 2. Section 3 considers applying relative belief to the multiple testing problem. In Section 4 some practical applications are presented with special emphasis on regression problems including the situation where the number of predictors exceeds the number of observations.

There have been a number of priors proposed for the sparsity problem in the literature. A common choice is the spike-and-slab prior due to George and McCulloch (1993). Other priors are available that are known to produce sparsity at least in connection with MAP estimation. For example, the Laplace prior has a close connection with sparsity via the LASSO, as discussed in Tibshirani (1996), Park and Casella (2008) and Hastie, Tibshirani and Wainwright (2015). The horseshoe prior of Carvalho, Polson and Scott (2009) will also produce

sparsity through estimation. Both the Laplace and the horseshoe prior can be checked for prior-data conflict but the check is not explicitly for sparsity, only that the location and scaling of the prior is appropriate. This reinforces a comment in Carvalho, Polson and Scott (2009) that the use of the spike-and-slab prior represents a kind of "gold standard" for sparsity. The problem with the spike-and-slab prior, and not possessed to the same degree by the Laplace or horseshoe prior, is the difficulty of the computations when implementing the posterior analysis. Basically there are 2^k mixture components to the posterior and this becomes computationally intractable as k rises. Various approaches can be taken to try and deal with this issue such as those developed in Rockova and George (2014). Any of these priors can be used with the approach taken here but our methodology does not require that a prior induces sparsity in any sense. Indeed sparsity is induced only when the evidence, as measured by how the data changes beliefs, points to this. This is a strong point of the approach as sparsity arises only through an assessment of the evidence.

2 Inferences Based on Relative Belief Ratios

Suppose now that, in addition to the statistical model $\{f_\theta : \theta \in \Theta\}$ there is a prior π on Θ. After observing data x, the posterior distribution of θ is then given by the density $\pi(\theta \,|\, x) = \pi(\theta) f_\theta(x)/m(x)$ where $m(x) = \int_\Theta \pi(\theta) f_\theta(x)\, d\theta$ is the prior predictive density of x. For a parameter of interest $\Psi(\theta) = \psi$ denote the prior and posterior densities of ψ by π_Ψ and $\pi_\Psi(\cdot \,|\, x)$, respectively. The *relative belief ratio* for a value ψ is then defined by $RB_\Psi(\psi \,|\, x) = \lim_{\delta \to 0} \Pi_\Psi(N_\delta(\psi) \,|\, x)/\Pi_\Psi(N_\delta(\psi))$ where $N_\delta(\psi)$ is a sequence of neighborhoods of ψ converging (nicely) to $\{\psi\}$ as $\delta \to 0$. Quite generally

$$RB_\Psi(\psi \,|\, x) = \pi_\Psi(\psi \,|\, x)/\pi_\Psi(\psi), \tag{1}$$

the ratio of the posterior density to the prior density at ψ. So $RB_\Psi(\psi \,|\, x)$ is measuring how beliefs have changed concerning ψ being the true value from *a priori* to *a posteriori*. When $RB_\Psi(\psi \,|\, x) > 1$ the data have lead to an increase in the probability that ψ is correct and so there is evidence in favor of ψ, when $RB_\Psi(\psi \,|\, x) < 1$ the data have lead to a decrease in the probability that ψ is correct and so there is evidence against ψ, and when $RB_\Psi(\psi \,|\, x) = 1$ there is no evidence either way. While there are numerous quantities that measure evidence through change in belief, the relative belief ratio is perhaps the simplest and it possesses many nice properties as discussed in Evans (2015). For example, $RB_\Psi(\psi \,|\, x)$ is invariant under smooth changes of variable and also invariant to the choice of the support measure for the densities. As such all relative belief inferences possess this invariance which is not the case for many Bayesian inferences such as using a posterior mode (MAP) or expectation for estimation.

The value $RB_\Psi(\psi_0 \,|\, x)$ measures the evidence for the hypothesis $H_0 = \{\theta : \Psi(\theta) = \psi_0\}$. It is also necessary to calibrate whether this is strong or weak evidence for or against H_0. Certainly the bigger $RB_\Psi(\psi_0 \,|\, x)$ is than 1, the more evidence there is in favor of ψ_0 while the smaller $RB_\Psi(\psi_0 \,|\, x)$ is than 1, the more

evidence there is against ψ_0. But what exactly does a value of $RB_\Psi(\psi_0 \mid x) = 20$ mean? It would appear to be strong evidence in favor of ψ_0 because beliefs have increased by a factor of 20 after seeing the data. But what if other values of ψ had even larger increases? A calibration of (1) is then given by the *strength*

$$\Pi_\Psi(RB_\Psi(\psi \mid x) \leq RB_\Psi(\psi_0 \mid x) \mid x), \tag{2}$$

namely, the posterior probability that the true value of ψ has a relative belief ratio no greater than that of the hypothesized value ψ_0. While (2) may look like a p-value, it has a very different interpretation. For when $RB_\Psi(\psi_0 \mid x) < 1$, so there is evidence against ψ_0, then a small value for (2) indicates a large posterior probability that the true value has a relative belief ratio greater than $RB_\Psi(\psi_0 \mid x)$ and so there is strong evidence against ψ_0 while only weak evidence against if (2) is big. If $RB_\Psi(\psi_0 \mid x) > 1$, so there is evidence in favor of ψ_0, then a large value for (2) indicates a small posterior probability that the true value has a relative belief ratio greater than $RB_\Psi(\psi_0 \mid x))$ and so there is strong evidence in favor of ψ_0, while a small value of (2) only indicates weak evidence in favor of ψ_0. Notice that in $\{\psi : RB_\Psi(\psi \mid x) \leq RB_\Psi(\psi_0 \mid x)\}$, the 'best' estimate of ψ is given by ψ_0 because the evidence for this value is the largest in the set.

Various results have been established in Baskurt and Evans (2013) and Evans (2015) supporting both (1), as the measure of the evidence, and (2) as a measure of the strength of that evidence. For example, the following simple inequalities, see Baskurt and Evans (2013), are useful in assessing the strength as computing (2) can be avoided, namely, $\Pi_\Psi(\{\psi_0\} \mid x) \leq \Pi_\Psi(RB_\Psi(\psi \mid x) \leq RB_\Psi(\psi_0 \mid x) \mid x) \leq RB_\Psi(\psi_0 \mid x)$. So if $RB_\Psi(\psi_0 \mid x) > 1$ and $\Pi_\Psi(\{\psi_0\} \mid x)$ is large, there is strong evidence in favor of ψ_0 while, if $RB_\Psi(\psi_0 \mid x) < 1$ is very small, then there is immediately strong evidence against ψ_0.

The use of $RB_\Psi(\psi_0 \mid x)$ to assess the hypothesis H_0 also possesses optimality properties. For example, let A be a subset of the sample space such that whenever $x \in A$, the hypothesis is accepted. If $M(A \mid H_0)$ denotes the prior predictive probability of A given that H_0 is true, then $M(A \mid H_0)$ is the prior probability of accepting H_0 when it is true. The relative belief acceptance region is naturally given by $A_{rb}(\psi_0) = \{x : RB_\Psi(\psi_0 \mid x) > 1\}$. Similarly, let R be a subset such that whenever $x \in R$, the hypothesis is rejected and let $M(R \mid H_0^c)$ denote the prior predictive probability of R given that H_0 is false. The relative belief rejection region is then naturally given by $R_{rb}(\psi_0) = \{x : RB_\Psi(\psi_0 \mid x) < 1\}$. The following result is proved in Evans (2015) as Proposition 4.7.9.

Theorem 1 *(i) The acceptance region $A_{rb}(\psi_0)$ minimizes $M(A)$ among all acceptance regions A satisfying $M(A \mid H_0) \geq M(A_{rb}(\psi_0) \mid H_0)$. (ii) The rejection region $R_{rb}(\psi_0)$ maximizes $M(R)$ among all rejection regions R satisfying $M(R \mid H_0) \leq M(R_{rb}(\psi_0) \mid H_0)$.*

To see the meaning of this result note that

$$\begin{aligned} &M(A) = E_{\Pi_\Psi}(M(A \mid \psi)) = E_{\Pi_\Psi}(I_{H_0^c}(\psi) M(A \mid \psi)) + \Pi_\Psi(H_0) M(A \mid H_0) \\ &\geq M(A_{rb}(\psi_0)) = E_{\Pi_\Psi}(I_{H_0^c}(\psi) M(A_{rb}(\psi_0) \mid \psi)) + \Pi_\Psi(H_0) M(A_{rb}(\psi_0) \mid H_0). \end{aligned}$$

Therefore, if $\Pi_\Psi(H_0) = 0$, then $A_{rb}(\psi_0)$ minimizes $E_{\Pi_\Psi}(I_{\{\psi_0\}^c}(\psi)M(A \mid \psi))$, the prior probability that H_0 is accepted given that it is false, among all acceptance regions A satisfying $M(A \mid \psi_0) \geq M(A_{rb}(\psi_0) \mid \psi_0)$ and when $\Pi_\Psi(H_0) > 0$, then $A_{rb}(\psi_0)$ minimizes this probability among all acceptance regions A satisfying $M(A \mid H_0) = M(A_{rb}(\psi_0) \mid H_0)$. Also,

$$M(R) = E_{\Pi_\Psi}(M(R \mid \psi)) = E_{\Pi_\Psi}(I_{H_0^c}(\psi)M(R \mid \psi)) + \Pi_\Psi(H_0)M(R \mid H_0)$$
$$\leq M(R_{rb}(\psi_0)) = E_{\Pi_\Psi}(I_{H_0^c}(\psi)M(R_{rb}(\psi_0) \mid \psi)) + \Pi_\Psi(H_0)M(R_{rb}(\psi_0) \mid H_0).$$

Therefore, if $\Pi_\Psi(\{\psi_0\}) = 0$, then $R_{rb}(\psi_0)$ maximizes $E_{\Pi_\Psi}(I_{\{\psi_0\}^c}(\psi)M(R \mid \psi))$, the prior probability that H_0 is rejected given that it is false, among all rejection regions R satisfying $M(R \mid H_0) \leq M(R_{rb}(\psi_0) \mid H_0)$ and when $\Pi_\Psi(H_0\}) > 0$, then $R_{rb}(\psi_0)$ maximizes this probability among all rejection regions R satisfying $M(R \mid \psi_0) = M(R_{rb}(\psi_0) \mid \psi_0)$.

Note that it does not make sense to accept or reject H_0 when $RB_\Psi(\psi_0 \mid x) = 1$. Also, under i.i.d. sampling, $M(A_{rb}(\psi_0) \mid H_0) \to 1$ and $M(R_{rb}(\psi_0) \mid H_0) \to 0$ as sample size increases, so these quantities can be controlled by design. When $\Pi_\Psi(\{\psi_0\}) = 0$, then $M(A_{rb}(\psi_0) \mid H_0) = 1 - M(R_{rb}(\psi_0) \mid H_0)$ and so controlling $M(A_{rb}(\psi_0) \mid H_0)$ is controlling the "size" of the test. In general, $E_{\Pi_\Psi}(I_{\{\psi_0\}^c}(\psi)M(R_{rb}(\psi_0) \mid \psi))$ can be thought of as the Bayesian power of the relative belief test. Note that it is reasonable to set either the probability of a false positive or the probability of a true negative as part of design and so the theorem is an optimality result with practical import. It is easily seen that the proof of Theorem 1 can be generalized to obtain optimality results for the acceptance region $A_{rb,q}(\psi_0) = \{x : RB_\Psi(\psi_0 \mid x) > q\}$ and for the rejection region $R_{rb,q}(\psi_0) = \{x : RB_\Psi(\psi_0 \mid x) < q\}$. The following inequality is useful in Section 3 in controlling error rates.

Theorem 2 $M\left(R_{rb,q}(\psi_0) \mid \psi_0\right) \leq q$.

Proof: By the Savage-Dickey result, see Proposition 4.2.7 in Evans (2015), $RB_\Psi(\psi_0 \mid x) = m(x \mid \psi_0)/m(x)$. Now $E_{M(\cdot \mid \psi_0)}(m(x)/m(x \mid \psi_0)) = 1$ and so by Markov's inequality, $M\left(R_{rb,q}(\psi_0) \mid \psi_0\right) = M\left(m(x)/m(x \mid \psi_0) > 1/q \mid \psi_0\right) \leq q$.

There is another issue associated with using $RB_\Psi(\psi_0 \mid x)$ to assess the evidence that ψ_0 is the true value. One of the key concerns with Bayesian inference methods is that the choice of the prior can bias the analysis in various ways. For example, in many problems Bayes factors and relative belief ratios can be made arbitrarily large by choosing the prior to be increasingly diffuse. This phenomenon is associated with the Jeffreys-Lindley paradox and clearly indicates that it is possible to bias results by the choice of the prior.

An approach to dealing with this bias is discussed in Baskurt and Evans (2013). For, given a measure of evidence, it is possible to measure *a priori* whether or not the chosen prior induces bias either for or against ψ_0. The bias against ψ_0 is given by

$$M(RB_\Psi(\psi_0 \mid x) \leq 1 \mid \psi_0) = 1 - M(A_{rb}(\psi_0) \mid \psi_0) \tag{3}$$

as this is the prior probability that evidence will not be obtained in favor of ψ_0 when ψ_0 is true. If (3) is large, subsequently reporting, after seeing the data, that there is evidence against ψ_0 is not convincing.

To measure the bias in favor of ψ_0 choose values $\psi_0' \neq \psi_0$ such that the difference between ψ_0 and ψ_0' represents the smallest difference of practical importance. Then compute

$$M_T\left(RB_\Psi(\psi_0 \mid x) \geq 1 \mid \psi_0'\right) = 1 - M(R_{rb}(\psi_0) \mid \psi_0') \tag{4}$$

as this is the prior probability that we will not obtain evidence against ψ_0 when ψ_0 is false. Note that (4) tends to decrease as ψ_0' moves further away from ψ_0. When (4) is large, subsequently reporting, after seeing the data, that there is evidence in favor of ψ_0, is not convincing. For a fixed prior, both (3) and (4) decrease with sample size and so, in design situations, they can be used to set sample size and so control bias. Notice that $M(A_{rb}(\psi_0) \mid \psi_0)$ can be considered as the sensitivity and $M(R_{rb}(\psi_0) \mid \psi_0')$ as the specificity of the relative belief hypothesis assessment. These issues are further discussed in Evans (2015).

As $RB_\Psi(\psi \mid x)$ measures the evidence that ψ is the true value, it can also be used to estimate ψ. For example, the best estimate of ψ is clearly the value for which the evidence is greatest, namely, $\psi(x) = \arg\sup RB_\Psi(\psi \mid x)$. Associated with this is a γ-credible region $C_{\Psi,\gamma}(x) = \{\psi : RB_\Psi(\psi \mid x) \geq c_{\Psi,\gamma}(x)\}$ where $c_{\Psi,\gamma}(x) = \inf\{k : \Pi_\Psi(RB_\Psi(\psi \mid x) > k \mid x) \leq \gamma\}$. Notice that $\psi(x) \in C_{\Psi,\gamma}(x)$ for every $\gamma \in [0,1]$ and so, for selected γ, we can take the "size" of $C_{\Psi,\gamma}(x)$ as a measure of the accuracy of the estimate $\psi(x)$. The interpretation of $RB_\Psi(\psi \mid x)$ as the evidence for ψ forces the sets $C_{\Psi,\gamma}(x)$ to be the credible regions. For if ψ_1 is in such a region and $RB_\Psi(\psi_2 \mid x) \geq RB_\Psi(\psi_1 \mid x)$, then ψ_2 must be in the region as well as there is at least as much evidence for ψ_2 as for ψ_1. A variety of optimality results have been established for $\psi(x)$ and $C_{\Psi,\gamma}(x)$, see Evans (2015).

The view is taken here that any time continuous probability is used, then this is an approximation to a finite, discrete context. For example, if ψ is a mean and the response measurements are to the nearest centimeter, then of course the true value of ψ cannot be known to an accuracy greater than $1/2$ of a centimeter, no matter how large a sample we take. Furthermore, there are implicit bounds associated with any measurement process. As such the restriction is made here to discretized parameters that take only finitely many values. So when ψ is a continuous, real-valued parameter, it is discretized to the intervals $\ldots, (\psi_0 - 3\delta, \psi_0 - \delta], (\psi_0 - \delta, \psi_0 + \delta], (\psi_0 + \delta, \psi_0 + 3\delta], \ldots$ for some choice of $\delta > 0$, and there are only finitely may such intervals covering the range of possible values. It is of course possible to allow the intervals to vary in length as well. With this discretization, then we can take $H_0 = (\psi_0 - \delta, \psi_0 + \delta]$.

3 Inferences for Multiple Tests

Consider now the multiple testing problem discussed in Section 1. Let $\xi = \Xi(\theta) = \frac{1}{k}\sum_{i=1}^k I_{H_{0i}}(\Psi_i(\theta))$ be the proportion of the hypotheses H_{0i} that are true and suppose that Ψ_i is finite for each i, perhaps arising via a discretization

as discussed in Section 2. Note that the discreteness is essential to realistically determine what proportion of the hypotheses are correct, otherwise, under a continuous prior on Ψ, the prior distribution of $\Xi(\theta)$ is degenerate at 0. In an application it is desirable to make inference about the true value of $\xi \in \Xi = \{0, 1/k, 2/k, \ldots, 1\}$. This is based on the relative belief ratio $RB_\Xi(\xi \mid x) = \Pi(\Xi(\theta) = \xi \mid x)/\Pi(\Xi(\theta) = \xi)$. The appropriate estimate of ξ is then the relative belief estimate of Ξ, namely, $\xi(x) = \arg\sup_\xi RB_\Xi(\xi \mid x)$. The accuracy of this estimate is assessed using the size of $C_{\Xi,\gamma}(x)$ for some choice of $\gamma \in [0, 1]$. Also, hypotheses such as $H_0 = \{\theta : \Xi(\theta) \in [\xi_0, \xi_1]\}$, namely, the proportion true is at least ξ_0 and no greater than ξ_1, can be assessed using the relative belief ratio $RB(H_0 \mid x) = \Pi(\xi_0 \leq \Xi(\theta) \leq \xi_1 \mid x)/\Pi(\xi_0 \leq \Xi(\theta) \leq \xi_1)$ which equals $RB_\Xi(\xi_0 \mid x)$ when $\xi_0 = \xi_1$. The strength of this evidence can be assessed as previously discussed.

The estimate $\xi(x)$ can be used to control how many hypotheses are potentially accepted. For this select $k\xi(x)$ of the H_{0i} as being true from among those for which $RB_{\Psi_i}(\psi_{0i} \mid x) > 1$. Note that it does not make sense to accept H_{0i} as being true when $RB_{\Psi_i}(\psi_{0i} \mid x) < 1$ as there is evidence against this hypothesis. So, if there are fewer than $k\xi(x)$ satisfying $RB_{\Psi_i}(\psi_{0i} \mid x) > 1$, then fewer than this number should be accepted. If there are more than $k\xi(x)$ of the relative belief ratios satisfying $RB_{\Psi_i}(\psi_{0i} \mid x) > 1$, then some method will have to be used to select the $k\xi(x)$ which are potentially accepted. It is clear, however, that the logical way to do this is to order the H_{0i}, for which $RB_{\Psi_i}(\psi_{0i} \mid x) > 1$, based on their strengths $\Pi_\Psi(RB_{\Psi_i}(\psi_{0i} \mid x) \leq RB_{\Psi_i}(\psi_{0i} \mid x) \mid x)$, from largest to smallest, and accept at most the $k\xi(x)$ for which the evidence is strongest.

Note too that, if some of these strengths are indeed weak, there is no need to necessarily accept these hypotheses. The ultimate decision as to whether or not to accept a hypothesis is application dependent and is not statistical in nature. The role of statistics is to supply a clear statement of the evidence and its strength, while other criteria come into play when decisions are made. In any case, it is seen that inference about ξ is being used to help control the number of hypotheses accepted.

If, as is more common, control is desired of the number of false positives, then the relevant parameter of interest is $\upsilon = \Upsilon(\theta) = 1 - \Xi(\theta)$, the proportion of false hypotheses. Note that $\Pi(\Upsilon(\theta) = \upsilon) = \Pi(\Xi(\theta) = 1 - \upsilon)$ and so the relative belief estimate of υ satisfies $\upsilon(x) = 1 - \xi(x)$. Following the same procedure, the H_{0i} with $RB_{\Psi_i}(\psi_{0i} \mid x) < 1$ are then ranked via their strengths and at most $k\upsilon(x)$ are rejected.

The consistency of the procedure just described, for correctly identifying the H_{0i} that are true and those that are false, follows from results proved in Section 4.7.1 of Evans (2015) under i.i.d. sampling. In other words, as the amount of data increases, $\xi(x)$ converges to the proportion of H_{0i} that are true, each $RB(\psi_{0i} \mid x)$ converges to the largest possible value (always bigger than 1) when H_{0i} is true and converges to 0 when H_{0i} is false, and the evidence in favor or against converges to the strongest possible, depending on whether the hypothesis in question is true or false.

We refer to this procedure as the *multiple testing algorithm.* Consider first

a somewhat artificial example where many of the computations are easy but which fully demonstrates the relevant characteristics of the methodology.

Example 1. *Location normal.*

Suppose that there are k independent samples x_{ij} for $1 \leq i \leq k, 1 \leq j \leq n$ where the i-th sample is from a $N(\mu_i, \sigma^2)$ distribution with μ_i unknown and σ^2 known. It also assumed that prior knowledge about all the unknown μ_i is reflected in the statement that the $\mu_1, \ldots, \mu_k$ are i.i.d. from a $N(\mu_0, \lambda_0^2\sigma^2)$ distribution. It is easy to modify our development to allow the sample sizes to vary and to use a general multivariate normal prior, while the case where σ^2 is unknown is considered in Section 4. This context is relevant to the analysis of microarray data.

The value of (μ_0, λ_0^2) is determined via elicitation. For this it is supposed that it is known with virtual certainty that each $\mu_i \in (m_l, m_u)$ for specified values $m_l \leq m_u$. Here virtual certainty is interpreted to mean that the prior probability of this interval is at least 0.99 and other choices could be made. It is also supposed that $\mu_0 = (m_l + m_u)/2$. This implies that $\lambda_0 = (m_u - m_l)/(2\sigma\Phi^{-1}((1 + 0.99)/2))$. Following Evans and Jang (2011b), increasing the value of λ_0 implies a more weakly informative prior in this context and, as such, decreases the possibility of prior-data conflict.

The posterior distributions of the μ_i are then independent with $\mu_i \,|\, x \sim N(\mu_i(x), (n\lambda_0^2+1)^{-1}\lambda_0^2\sigma^2)$ where $\mu_i(x) = (n+1/\lambda_0^2)^{-1}(n\bar{x}_i+\mu_0/\lambda_0^2)$. Given that the measurements are taken to finite accuracy, it is not necessarily realistic to test $\mu_i = \mu_0$. As such, a value $\delta > 0$ is specified so that $H_{0i} = (\mu_0-\delta/2, \mu_0+\delta/2]$ in a discretization of the parameter space into a finite number of intervals, each of length δ, as well as two tail intervals. Then for $T \in \mathbb{N}$ there are $2T+1$ intervals of the form $(\mu_0 + (t - 1/2)\delta, \mu_0 + (t + 1/2)\delta]$, for $t \in \{-T, -T + 1, \ldots, T\}$ that span (m_l, m_u), together with two additional tail intervals $(-\infty, \mu_0 - (T+1/2)\delta]$ and $(\mu_0 + (T + 1/2)\delta, \infty)$ to cover the full range. The relative belief ratio for the t-th interval for μ_i is then given by

$$RB_i((\mu_0 + (t - 1/2)\delta, \mu_0 + (t + 1/2)\delta] \,|\, x)$$
$$= \frac{\left\{\begin{array}{c} \Phi((n\lambda_0^2 + 1)^{1/2}(\mu_0 + (t + 1/2)\delta - \mu_i(x))/\lambda_0\sigma) - \\ \Phi((n\lambda_0^2 + 1)^{1/2}(\mu_0 + (t - 1/2)\delta - \mu_i(x))/\lambda_0\sigma) \end{array}\right\}}{\Phi((t + 1/2)\delta/\lambda_0\sigma) - \Phi((t - 1/2)\delta/\lambda_0\sigma)} \tag{5}$$

with a similar formula for the tail intervals. When δ is small this is approximated by the ratio of the posterior to prior densities of μ_i evaluated at $\mu_0 + t\delta$. Then $RB(H_{0i} \,|\, x) = RB_i((\mu_0-\delta/2, \mu_0+\delta/2] \,|\, x)$ gives the evidence for or against H_{0i} and the strength of this evidence is computed using the discretized posterior distribution. Notice that (5) converges to ∞ as $\lambda_0 \to \infty$ and this is characteristic of other measures of evidence such as Bayes factors. As discussed in Evans (2015), this is one of the reasons why calibrating (1) via (2) is necessary.

As previously noted, it is important to take account of the bias in the prior. To simplify matters the continuous approximation is used here as this makes little difference for the discussion of bias concerning inference about μ_i (see

Tables 2 and 3). The bias against $\mu_i = \mu_0$ equals

$$M(RB_i(\mu_0 \,|\, x) \leq 1 \,|\, \mu_0) = 2(1 - \Phi(a_n(1))) \tag{6}$$

where

$$a_n(q) = \begin{cases} \{(1 + 1/n\lambda_0^2)\log((n\lambda_0^2 + 1)/q^2)\}^{1/2}, & q^2 \leq n\lambda_0^2 + 1 \\ 0, & q^2 > n\lambda_0^2 + 1. \end{cases}$$

Note that (6) converges to $2(1 - \Phi(1)) = 0.32$ as $\lambda_0 \to 0$ and to 0 as $\lambda_0 \to \infty$ and, for fixed λ_0, converges to 0 as $n \to \infty$. So there is never strong bias against $\mu_i = \mu_0$ and this is as expected since the prior is centered on μ_0. The bias in favor of $\mu_i = \mu_0$ is measured by

$$M(RB_i(\mu_0 \,|\, x) \geq 1 \,|\, \mu_0 \pm \delta/2) = \Phi(\sqrt{n}\delta/2\sigma + a_n(1)) - \Phi(\sqrt{n}\delta/2\sigma - a_n(1)). \tag{7}$$

As $\lambda_0 \to \infty$ then (7) converges to 1 so there is bias in favor of $\mu_i = \mu_0$ and this reflects what was obtained for the limiting value of (5). Also this decreases with increasing δ and goes to 0 as $n \to \infty$. So indeed bias of both types can be controlled by sample size. Perhaps the most important take away from this discussion, however, is that by using a supposedly noninformative prior with λ_0 large, bias in favor of the H_{0i} is being induced.

Consider first a simulated data set x when $k = 10, n = 5, \sigma = 1, \delta = 1, \mu_0 = 0, (m_l, m_u) = (-5, 5)$, so that $\lambda_0 = 10/2\Phi^{-1}(0.995) = 1.94$ and suppose $\mu_1 = \mu_2 = \cdots = \mu_7 = 0$ with the remaining $\mu_i = 2$. The relative belief ratio function $RB_\Xi(\cdot \mid x)$ is plotted in Figure 1. In this case the relative belief estimate $\xi(x) = 0.70$ is exactly correct. Table 1 gives the values of the $RB_i(0 \,|\, x)$ together with their strengths. It is clear that the multiple testing algorithm leads to 0 false positives and 0 false negatives. So the algorithm works perfectly on this data but of course it can't be expected to do as well when the three nonzero means move closer to 0. Also, it is worth noting that the strength of the evidence in favor of $\mu_i = 0$ is very strong for $i = 1, 2, 3, 5, 6, 7$ but only moderate when $i = 4$. The strength of the evidence against $\mu_i = 0$ is very strong for $i = 8, 9, 10$. Note that the maximum possible value of $RB_i((\mu_0 - \delta/2, \mu_0 + \delta/2] \,|\, x)$ here is $(2\Phi(\delta/2\lambda_0\sigma) - 1)^{-1} = 4.92$, so indeed some of the relative belief ratios are relatively large.

Now consider basically the same context but where $k = 1000$ and $\mu_1 = \cdots = \mu_{700} = 0$ while the remaining 300 satisfy $\mu_i = 2$. The relative belief ratio $RB_\Xi(\cdot \mid x)$ is plotted in Figure 2. In this case $\xi(x) = 0.47$ which is a serious underestimate. As such the multiple testing algorithm will not record enough acceptances and so will fail.

This problem arises due to the independence assumption on the μ_i. For the prior distribution of $k\Xi(\theta)$ is binomial$(k, 2\Phi(\delta/2\lambda_0\sigma) - 1)$ and the prior distribution of $k\Upsilon(\theta)$ is binomial$(k, 2(1 - \Phi(\delta/2\lambda_0\sigma)))$. So the *a priori* expected proportion of true hypotheses is $2\Phi(\delta/2\lambda_0\sigma) - 1$ and the expected proportion of false hypotheses is $2(1 - \Phi(\delta/2\lambda_0\sigma))$. When $\delta/2\lambda_0\sigma$ is small, as when the amount of sampling variability or the diffuseness of the prior are large, then the prior on

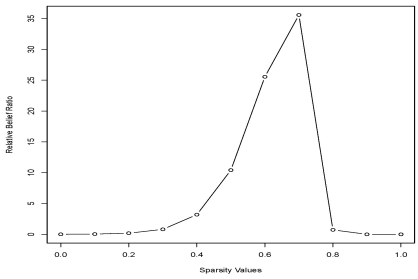

Figure 1: A plot of the relative belief ratio of Ξ when $n = 5, k = 10$ and 7 means equal 0 with the remaining means equal to 2 in Example 1.

i	1	2	3	4	5
μ_i	0	0	0	0	0
$RB_i(0 \mid x)$	3.27	3.65	2.98	1.67	3.57
Strength	1.00	1.00	1.00	0.37	1.00
i	6	7	8	9	10
μ_i	0	0	2	2	2
$RB_i(0 \mid x)$	3.00	3.43	2.09×10^{-4}	3.99×10^{-4}	8.80×10^{-3}
Strength	1.00	1.00	4.25×10^{-5}	8.11×10^{-5}	1.83×10^{-3}

Table 1: Relative belief ratios and strengths for the μ_i in Example 1 with $k = 10$.

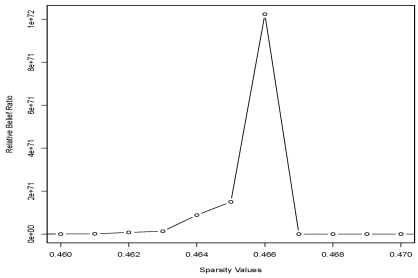

Figure 2: A plot of the relative belief ratio of Ξ when $n = 5, k = 1000$ and 700 means equal 0 with the remaining means equal to 2 in Example 1.

Ξ suggests a belief in many false hypotheses. When k is small, relatively small amounts of data can override this to produce accurate inferences about ξ or υ but otherwise large amounts of data are needed which may not be available. ■

Given that accurate inference about ξ and υ is often not feasible, we focus instead on protecting against too many false positives and false negatives. From the discussion in Example 1 it is seen that it is possible to produce bias in favor of the H_{0i} being true simply by using a diffuse prior. If our primary goal is to guard against too many false positives, then this biasing will work as we can make the prior conditional probabilities of the events $RB(\psi_{0i} \mid x) < 1$, given that H_{0i} is true, as small as desirable by choosing such a prior. This could be seen as a way of creating a higher bar for a positive result. The price we pay for this, however, is too many false negatives. As such we consider another way of "raising the bar". Given that $RB(\psi_{0i} \mid x)$ is measuring evidence, a natural approach is to simply choose constants $0 < q_R \leq 1 \leq q_A$ and classify H_{0i} as accepted when $RB(\psi_{0i} \mid x) > q_A$ and rejected when $RB(\psi_{0i} \mid x) < q_R$. The strengths can also be quoted to assess the reliability of these inferences. Provided q_R is greater than the minimum possible value of $RB(\psi_{0i} \mid x)$, and this is typically 0, and q_A is chosen less than the maximum possible value of $RB(\psi_{0i} \mid x)$, and this is 1 over the prior probability of H_{0i}, then this procedure is consistent as the amount of data increases. In fact, the related estimates of ξ and υ are also consistent. The price paid for this is that a hypothesis will not be classified whenever $q_R \leq RB(\psi_{0i} \mid x) \leq q_A$. Not classifying a hypothesis implies that there is not enough evidence for this purpose and more data is required. This approach is referred to as the *relative belief multiple testing algorithm.*

It remains to determine q_A and q_R. Consider first protecting against too many false positives. The *a priori* conditional prior probability, given that H_{0i} is true, of finding evidence against H_{0i} less than q_R is $M(RB_i(\psi_{0i} \mid X) < q_R \mid \psi_{0i}) = M(m(X \mid \psi_{0i})/m(x) < q_R \mid \psi_{0i}) \leq q_R$ where the final inequality follows from Theorem 2. Naturally, we want the probability of a false negative to be small and so choosing q_R small accomplishes this. The *a priori* probability that a randomly selected hypothesis produces a false positive is

$$\frac{1}{k}\sum_{i=1}^{k} M(RB_i(\psi_{0i} \mid X) < q_R \mid \psi_{0i}) \tag{8}$$

which by Theorem 2 is bounded above by q_R and so converges to 0 as $q_R \to 0$. Also, for fixed q_R, (8) converges to 0 as the amount of data increases. More generally q_R can be allowed to depend on i but when the ψ_i are similar in nature this does not seem necessary. Furthermore, it is not necessary to weight the hypotheses equally so a randomly chosen hypothesis with unequal probabilities could be relevant in certain circumstances. In any case, controlling the value of (8), whether by sample size or by the choice of q_R, is clearly controlling for false positives. Suppose there is proportion p_{FP} of false positives that is just tolerable in a problem. Then q_R can be chosen so that (8) is less than or equal to p_{FP} and note that $q_R = p_{FP}$ satisfies this.

Similarly, if $\psi'_{0i} \neq \psi_{0i}$ then $M(RB_i(\psi_{0i} \mid X) > q_A \mid \psi'_{0i})$ is the prior probability of finding evidence for H_{0i} when ψ'_{0i} is the true value. For a given effect size δ of practical importance it is natural to take $\psi'_{0i} = \psi_{0i} \pm \delta/2$. In typical applications this probability becomes smaller the "further" ψ'_{0i} is from ψ_{0i} and so choosing q_A to make this probability small will make it small for all alternatives. Under these circumstances the *a priori* probability that a randomly selected hypothesis produces a false negative is bounded above by

$$\frac{1}{k}\sum_{i=1}^{k} M(RB_i(\psi_{0i} \mid X) > q_A \mid \psi'_{0i}). \tag{9}$$

As $q_A \to \infty$, or as the amount of data increases with q_A fixed, then (9) converges to 0 so the number of false negatives can be controlled. If there is proportion p_{FN} of false negatives that is just tolerable in a problem, then q_A can be chosen so that (9) is less than or equal to p_{FN}.

The following optimality result holds for relative belief multiple testing.

Corollary 3 *(i) Among all procedures where the prior probability of accepting H_{0i}, when it is true, is at least $M(RB_i(\psi_{0i} \mid X) > q_A \mid \psi_{0i})$ for $i = 1, \ldots, k$, the relative belief multiple testing algorithm minimizes the prior probability that a randomly chosen hypothesis is accepted. (ii) Among all procedures where the prior probability of rejecting H_{0i}, when it is true, is less than or equal to $M(RB_i(\psi_{0i} \mid X) < q_R \mid \psi_{0i})$, then the relative belief multiple testing algorithm maximizes the prior probability that a randomly chosen hypothesis is rejected.*

Proof: For (i) consider a procedure for multiple testing and let A_i be the set of data values where H_{0i} is accepted. Then by hypothesis $M(RB_i(\psi_{0i} \mid X) > q_A \mid \psi_{0i}) \leq M(A_i \mid \psi_{0i})$ and by the analog of Theorem 1 $M(A_i) \geq M(RB_i(\psi_{0i} \mid X) > q_A)$. Applying this to a randomly chosen H_{0i} gives the result. The proof of (ii) is basically the same.

Applying the same discussion as after Theorem 1, it is seen that, under reasonable conditions, the relative belief multiple testing algorithm minimizes the prior probability of accepting a randomly chosen H_{0i} when it is false and maximizes the prior probability of rejecting a randomly chosen H_{0i} when it is false.

Consider now the application of the relative belief multiple testing algorithm in the previous example.

Example 2. *(Example 1 continued)*

In this context (8) equals $M(RB_i(\mu_0 \mid x) < q_R \mid \mu_0) = 2(1 - \Phi(a_n(q_R))$ for all i and so this is the value of (8). Therefore, q_R is chosen to make this number suitably small. Table 2 records values for (8). From this it is seen that for small n there can be some bias against H_{0i} ($q_R = 1$) and so the prior probability of obtaining false positives is perhaps too large. Table 2 demonstrates that choosing a smaller value of q_R can adequately control the prior probability of false positives.

n	λ_0	q_R	(8)	n	λ_0	q_R	(8)
1	1	1	0.239 (0.228)	5	1	1	0.143 (0.097)
		1/2	0.041 (0.030)			1/2	0.051 (0.022)
		1/10	0.001 (0.000			1/10	0.006 (0.001)
	2	1	0.156 (0.146)		2	1	0.074 (0.041)
		1/2	0.053 (0.045)			1/2	0.031 (0.013)
		1/10	0.005 (0.004)			1/10	0.005 (0.001)
	10	1	0.031 (0.026)		10	1	0.013 (0.004)
		1/2	0.014 (0.011)			1/2	0.006 (0.002)
		1/10	0.002 (0.002)			1/10	0.001 (0.001)

Table 2: Prior probability a randomly chosen hypothesis produces a false positive, continuous and discretized ( ) versions, in Example 2.

n	λ_0	q_A	(9)	n	λ_0	q_A	(9)
1	1	1.0	0.704 (0.715)	5	1	1.0	0.631 (0.702)
		1.2	0.527 (0.503)			2.0	0.302 (0.112)
		1.4	0.141 (0.000)			2.4	0.095 (0.000)
	2	1.0	0.793 (0.805)		2	1.0	0.747 (0.822)
		2.0	0.359 (0.304)			3.0	0.411 (0.380)
		2.2	0.141 (0.000)			4.5	0.084 (0.000)
	10	1.0	0.948 (0.955)		10	1.0	0.916 (0.961)
		5.0	0.708 (0.713)			10.0	0.552 (0.588)
		10.0	0.070 (0.000)			22.0	0.080 (0.000)

Table 3: Prior probability a randomly chosen hypothesis produces a false negative when $\delta/\sigma = 1$, continuous and discretized ( ) versions, in Example 2.

For false negatives consider (9) where

$$M(RB_i(\mu_0 \,|\, x) > q_A \,|\, \mu_0 \pm \delta/2) = \begin{cases} \Phi(\sqrt{n}\delta/2\sigma + a_n(q_A)) - & 1 \le q_A^2 \le \\ \Phi(\sqrt{n}\delta/2\sigma - a_n(q_A)), & n\lambda_0^2 + 1 \\ 0, & q_A^2 > n\lambda_0^2 + 1. \end{cases}$$

for all i. It is easy to show that this is monotone decreasing in δ and so it is an upper bound on the expected proportion of false negatives among those hypotheses that are actually false. The cutoff q_A can be chosen to make this number as small as desired. When $\delta/\sigma \to \infty$, then (9) converges to 0 and increases to $2\Phi(a_n(q_A)) - 1$ as $\delta/\sigma \to 0$. Table 3 records values for (9) when $\delta/\sigma = 1$ so that the μ_i differ from μ_0 by one half of a standard deviation. There is clearly some improvement but still the biasing in favor of false negatives is readily apparent. It would seem that taking $q_A = \sqrt{n\lambda_0^2 + 1}$ gives the best results but this could be considered as quite conservative. It is also worth remarking that all the entries in Table 3 can be considered as very conservative when large effect sizes are expected.

Now consider the situation that led to Figure 2. For this $k = 1000, n = 5$

Decision	$\mu = 0$	$\mu = 2$
Accept $\mu = 0$ using $q_A = 1.0$	666	3
Reject $\mu = 0$ using $q_R = 1.0$	34	297
Not classified	0	0
Accept $\mu = 0$ using $q_A = 3.0$	419	0
Reject $\mu = 0$ using $q_R = 0.5$	9	287
Not classified	272	13

Table 4: Confusion matrices for Example 2 with $k = 1000$ when 700 of the μ_i equal 0 and 300 equal 2.

and $\lambda_0 = 1.94$ is the elicited value. From Table 2 with $q_R = 1.0$, about 8% false positives are expected *a priori* and from Table 3 with $q_A = 1.0$, a worst case upper bound on the *a priori* expected percentage of false negatives is about 75%. The top part of Table 4 indicates that with $q_R = q_A = 1.0$, then 4.9% (34 of 700) false positives and 0.1% (3 of 300) false negatives were obtained. With these choices of the cutoffs all hypotheses are classified. Certainly the upper bound 75% seems far too pessimistic in light of the results, but recall that Table 3 is computed at the false values $\mu = \pm 0.5$. The relevant *a priori* expected percentage of false negatives when $\mu = \pm 2.0$ is about 3.5%. The bottom part of Table 4 gives the relevant values when $q_R = 0.5$ and $q_A = 3.0$. In this case there are 2.1% (9 of 428) false positives and 0% false negatives but 39.9% (272 out of 700) of the true hypotheses and 4.3% (13 out 300) of the false hypotheses were not classified as the relevant relative belief ratio lay between q_R and q_A. So in this case being more conservative has reduced the error rates with the price being a large proportion of the true hypotheses don't get classified. The procedure has worked well in this example but of course the error rates can be expected to rise when the false values move towards the null and improve when they move away from the null. ■

What is implemented in an application depends on the goals. If the primary purpose is to protect against false positives, then Table 2 indicates that this is accomplished fairly easily. Protecting against false negatives is more difficult. Since the actual effect sizes are not known a decision has to be made. Note that choosing a cutoff is equivalent to saying that one will only accept H_{0i} if the belief in the truth of H_{0i} has increased by a factor at least as large as q_A. Computations such as in Table 3 can be used to provide guidance but there is no avoiding the need to be clear about what effect sizes are deemed to be important or the need to obtain more data when this is necessary. One comforting aspect of the methodology is that error rates are effectively controlled but there may be many true hypotheses not classified.

The idea of controlling the prior probability of a randomly chosen hypothesis yielding a false positive or a false negative via (8) or (9), respectively, can be extended. For example, consider the prior probability that a random sample of

l from k hypotheses yields at least one false positive

$$\frac{1}{\binom{k}{l}} \sum_{\{i_1,\ldots,i_l\}\subset\{1,\ldots,k\}} M\left(\begin{array}{c} \text{at least one of } RB_{i_j}(\psi_{0i_j} \,|\, X) < q_R \\ \text{for } j = 1,\ldots,l \,|\, \psi_{0i_1},\ldots,\psi_{0i_l} \end{array} \right). \tag{10}$$

In the context of the examples of this paper, and many others, the term in (10) corresponding to $\{i_1,\ldots,i_l\}$ equals $M(\text{at least one of } RB_{i_j}(\psi_{0i_j} \,|\, X) < q_R \text{ for } j = 1,\ldots,l \,|\, \psi_0)$. The following result, whose proof is given in the Appendix, then leads to an interesting property for (10).

Lemma 4 *For probability model $(\Omega, \mathcal{F}, P)$, the probability that at least one of $l \leq k$ randomly selected events from $\{A_1,\ldots,A_k\} \subset \mathcal{F}$ occurs is an upper bound on the probability that at least one of $l' \leq l$ randomly selected events from $\{A_1,\ldots,A_k\} \subset \mathcal{F}$ occur.*

It then follows, by taking $A_i = \{x : RB_i(\psi_{0i} \,|\, x) < q_R\}$, that (10) is an upper bound on the prior probability that a random sample of l' hypotheses yields at least one false positive whenever $l' \leq l$. So (10) leads to a more rigorous control over the possibility false positives, if so desired. A similar result is obtained for false negatives.

4 Applications

One application of the relative belief multiple testing algorithm is to the problem of inducing sparsity.

Example 3. *Testing for sparsity.*

The position taken here is that, rather than trying to induce sparsity through an estimation procedure, an assessment of sparsity is made through an explicit measure of the evidence as to whether or not a particular assignment is valid. Virtues of this approach are that it is based on a measure of evidence and it is not dependent on the form of the prior as any prior can be used. Also, it has some additional advantages when prior-data conflict is taken into account.

Consider the context of Example 1. A natural approach to inducing sparsity is to estimate μ_i by μ_0 whenever $RB_i(\mu_0 \,|\, x) > q_A$. From the simulation it is seen that this works extremely well when $q_A = 1$ for both $k = 10$ and $k = 1000$. It also works when $k = 1000$ and $q_A = 3$, in the sense that the error rate is low, but it is conservative in the amount of sparsity it induces in that case. Again the goals of the application will dictate what is appropriate.

A common method for inducing sparsity is to use a penalized estimator as in the LASSO introduced in Tibshirani (1996). It is noted there that the LASSO is equivalent to MAP estimation when using a product of independent Laplace priors. This aspect was pursued further in Park and Casella (2006) which adopted a more formal Bayesian approach but still used MAP.

Consider then a product of independent Laplace priors for the prior on μ, namely, $(\sqrt{2}\lambda_0\sigma)^{-k} \exp\{-(\sqrt{2}/\lambda_0\sigma)\sum_{i=1}^k |\mu_i - \mu_0|\}$ where σ is assumed known

Decision	$\mu = 0$	$\mu = 2$
Accept $\mu = 0$ using $q_A = 1.0$	227	0
Reject $\mu = 0$ using $q_A = 1.0$	473	300

Table 5: Confusion matrices using LASSO with $k = 1000$ when 700 of the μ_i equal 0 and 300 equal 2 in Example 3

and μ_0, λ_0 are hyperparameters. Note that each Laplace prior has mean μ_0 and variance $\lambda_0^2\sigma^2$. Using the elicitation algorithm provided in Example 1, but replacing the normal prior with a Laplace prior, leads to the assignment $\mu_0 = (m_l + m_u)/2$, $\lambda_0 = (m_u - m_l)/2\sigma G^{-1}(0.995)$ where $G^{-1}(p) = 2^{-1/2}\log 2p$ when $p \leq 1/2, G^{-1}(p) = -2^{-1/2}\log 2(1-p)$ when $p \geq 1/2$ and G^{-1} denotes the quantile function of a Laplace distribution with mean 0 and variance 1. With the specifications used in the simulations of Example 1, this leads to $\mu_0 = 0$ and $\lambda_0 = 1.54$ which implies a smaller variance than the value $\lambda_0 = 1.94$ used with the normal prior and so the Laplace prior is more concentrated about 0.

The posteriors for the μ_i are independent with the density for μ_i proportional to $\exp\{-n(\bar{x}_i - \mu_i)^2/2\sigma^2 - \sqrt{2}|\mu_i - \mu_0|/\lambda_0\sigma\}$ giving the MAP estimator

$$\mu_{i\text{MAP}}(x) = \begin{cases} \bar{x}_i + \sqrt{2}\sigma/\lambda_0 n, & \bar{x}_i < \mu_0 - \sqrt{2}\sigma/\lambda_0 n \\ \mu_0, & \mu_0 - \sqrt{2}\sigma/\lambda_0 n \leq \bar{x}_i \leq \mu_0 + \sqrt{2}\sigma/\lambda_0 n \\ \bar{x}_i - \sqrt{2}\sigma/\lambda_0 n, & \bar{x}_i > \mu_0 + \sqrt{2}\sigma/\lambda_0 n. \end{cases}$$

The MAP estimate of μ_i is sometimes forced to equal μ_0 although this effect is negligible whenever $\sqrt{2}\sigma/\lambda_0 n$ is small.

The LASSO induces sparsity through estimation by taking λ_0 to be small. By contrast the evidential approach, based on the normal prior and the relative belief ratio, induces sparsity through taking λ_0 large. The advantage to this latter approach is that by taking λ_0 large, prior-data conflict is avoided. When taking λ_0 small, the potential for prior-data conflict rises as the true values can be deep into the tails of the prior. For example, for the simulations of Example 1 $\sqrt{2}\sigma/\lambda_0 n = 0.183$ which is smaller than the $\delta/2 = 0.5$ used in the relative belief approach with the normal prior. So it can be expected that the LASSO will do worse here and this is reflected in Table 5 where there are far too many false negatives. To improve this the value of λ_0 needs to be reduced although note that this is determined by an elicitation and there is the risk of then encountering prior-data conflict. Another possibility is to implement the evidential approach with the elicited Laplace prior and the discretization as done with the normal prior and then we can expect results similar to those obtained in Example 1.

It is also interesting to compare the MAP estimation approach and the relative belief approach with respect to the conditional prior probabilities of μ_i being assigned the value μ_0 when the true value actually is μ_0. It is easily seen that, based on the Laplace prior, $M(\mu_{i\text{MAP}}(x) = \mu_0 \,|\, \mu_0) = 2\Phi(\sqrt{2}/\lambda_0\sqrt{n}) - 1$ and this converges to 0 as $n \to \infty$ or $\lambda_0 \to \infty$. For the relative belief approach $M(RB_i(\mu_0 \,|\, x) > q_A \,|\, \mu_0)$ is the relevant probability. With either the normal or

Laplace prior $M(RB_i(\mu_0 \,|\, x) > q_A \,|\, \mu_0)$ converges to 1 both as $n \to \infty$ and as $\lambda_0 \to \infty$. In particular, with enough data the correct assignment is always made using relative belief but not with MAP based on the Laplace prior.

While the Laplace and normal priors work equally with the relative belief multiple testing algorithm, there don't appear to be any advantages to using the Laplace prior. One could argue too that the singularity of the Laplace prior at its mode makes it an odd choice and there doesn't seem to be a good justification for this. Furthermore, the computations are harder with the Laplace prior, particularly with more complex models. So using a normal prior seems preferable overall. ■

An example with considerable practical significance is now considered.

Example 4. *Full rank regression.*

Suppose the basic model is given by $y = \beta_0 + \beta_1 x_1 + \cdots + \beta_k x_k + z = \beta_0 + x'\beta_{1:k} + z$ where the x_i are predictor variables, $z \sim N(0, \sigma^2)$ and the β_i and σ^2 are unknown. The main issue in this problem is testing $H_{0i} : \beta_i = 0$ for $i = 1, \ldots, k$ to establish which variables have any effect on the response. The prior distribution of (β, σ^2) is taken to be

$$\beta \,|\, \sigma^2 \sim N_{k+1}(0, \sigma^2 \Sigma_0), 1/\sigma^2 \sim \text{gamma}_{\text{rate}}(\alpha_1, \alpha_2), \tag{11}$$

for some hyperparameters Σ_0 and (α_1, α_2). Note that this may entail subtracting a known, fixed constant from each y value so that the prior for β_0 is centered at 0. Taking 0 as the central value for the priors on the remaining β_i seems appropriate when the primary concern is whether or not each x_i is having any effect. Also, it will be assumed that the observed values of the predictor variable have been standardized so that for observations $(y, X) \in R^n \times R^{n\times(k+1)}$, where $X = (\mathbf{1}, \mathbf{x}_1, \ldots, \mathbf{x}_k)$, then $\mathbf{1}'\mathbf{x}_i = 0$ and $||\mathbf{x}_i||^2 = 1$ for $i = 1, \ldots, k$. The marginal prior for β_i is then $\{(\alpha_2/\alpha_1)\sigma_{0ii}^2\}^{1/2} t_{2\alpha_1}$ where $t_{2\alpha_1}$ denotes the t distribution on $2\alpha_1$ degrees of freedom, for $i = 0, \ldots, k$. Hereafter, we will take $\Sigma_0 = \lambda_0^2 I_{k+1}$ although it is easy to generalize to more complicated choices.

The elicitation of the hyperparameters is carried out via an extension of a method developed in Cao, Evans and Guttman (2014) for the multivariate normal distribution. Suppose that it is known with virtual certainty, based on our knowledge of the measurements being taken, that $\beta_0 + x'\beta_{1:k}$ will lie in the interval $(-m_0, m_0)$ for some $m_0 > 0$ for all $x \in R$ where R is a compact set centered at 0. On account of the standardization, $R \subset [-1, 1]^k$. Again 'virtual certainty' is interpreted as probability greater than or equal to γ where γ is some large probability like 0.99. Therefore, the prior on β must satisfy $2\Phi(m_0/\sigma\lambda_0\{1 + x'x\}^{1/2}) - 1 \geq \gamma$ for all $x \in R$ and this implies that

$$\sigma \leq m_0/\lambda_0 \tau_0 z_{(1+\gamma)/2} \tag{12}$$

where $\tau_0^2 = 1 + \max_{x\in R} ||x||^2 \leq 1 + k$ with equality when $R = [-1, 1]^k$.

An interval that will contain a response value y with virtual certainty, given predictor values x, is $\beta_0 + x'\beta_{1:k} \pm \sigma z_{(1+\gamma)/2}$. Suppose that we have lower and upper bounds s_1 and s_2 on the half-length of this interval so that

$s_1 \leq \sigma z_{(1+\gamma)/2} \leq s_2$ or, equivalently,

$$s_1/z_{(1+\gamma)/2} \leq \sigma \leq s_2/z_{(1+\gamma)/2} \tag{13}$$

holds with virtual certainty. Combining (13) with (12) implies $\lambda_0 = m_0/s_2\tau_0$.

To obtain the relevant values of α_1 and α_2 let $G(\alpha_1, \alpha_2, \cdot)$ denote the cdf of the $\text{gamma}_{\text{rate}}(\alpha_1, \alpha_2)$ distribution and note that $G(\alpha_1, \alpha_2, z) = G(\alpha_1, 1, \alpha_2 z)$. Therefore, the interval for $1/\sigma^2$ implied by (13) contains $1/\sigma^2$ with virtual certainty, when α_1, α_2 satisfy $G^{-1}(\alpha_1, \alpha_2, (1+\gamma)/2) = s_1^{-2} z^2_{(1+\gamma)/2}, G^{-1}(\alpha_1, \alpha_2, (1-\gamma)/2) = s_2^{-2} z^2_{(1-\gamma)/2}$, or equivalently

$$G(\alpha_1, 1, \alpha_2 s_1^{-2} z^2_{(1+\gamma)/2}) = (1+\gamma)/2, \tag{14}$$

$$G(\alpha_1, 1, \alpha_2 s_2^{-2} z^2_{(1-\gamma)/2}) = (1-\gamma)/2. \tag{15}$$

It is a simple matter to solve these equations for (α_1, α_2). For this choose an initial value for α_1 and, using (14), find z such that $G(\alpha_1, 1, z) = (1+\gamma)/2$, which implies $\alpha_2 = z/s_1^{-2} z^2_{(1+\gamma)/2}$. If the left-side of (15) is less (greater) than $(1-\gamma)/2$, then decrease (increase) the value of α_1 and repeat step 1. Continue iterating this process until satisfactory convergence is attained.

The methods discussed in Evans and Moshonov (2006) are available for checking the prior to see if it is contradicted by the data. Methods are specified there for checking each of the components in the hierarchy, namely, first check the prior on σ^2 and, if it passes, then check the prior on β. If conflict is found, then the methods discussed in Evans and Jang (2011b) are available to modify the prior appropriately.

Assuming that X is of rank $k+1$, the posterior of (β, σ^2) is given by

$$\begin{aligned} &\beta \,|\, y, \sigma^2 \sim N_{k+1}(\beta(X,y), \sigma^2 \Sigma(X)), \\ &1/\sigma^2 \,|\, y \sim \text{gamma}_{\text{rate}}((n+2\alpha_1)/2, \alpha_2(X,y)/2), \end{aligned} \tag{16}$$

where $b = (X'X)^{-1}X'y, \beta(X,y) = \Sigma(X)X'Xb, \Sigma(X) = (X'X + \Sigma_0^{-1})^{-1}$ and $\alpha_2(X,y) = ||y - Xb||^2 + (Xb)'(I_n - X\Sigma(X)X')Xb + 2\alpha_2$. Then the marginal posterior for β_i is given by $\beta_i(X,y) + \{\alpha_2(X,y)\sigma_{ii}(X)/(n+2\alpha_1)\}^{1/2} t_{n+2\alpha_1}$ and the relative belief ratio for β_i at 0 equals

$$\begin{aligned} RB_i(0 \,|\, X, y) = &\frac{\Gamma\left(\frac{n+2\alpha_1+1}{2}\right)\Gamma(\alpha_1)}{\Gamma\left(\frac{2\alpha_1+1}{2}\right)\Gamma\left(\frac{n+2\alpha_1}{2}\right)} \left(1 + \frac{\beta_i^2(X,y)}{\alpha_2(X,y)\sigma_{ii}(X)}\right)^{-\frac{n+2\alpha_1+1}{2}} \times \\ &\left(\frac{\alpha_2(X,y)\sigma_{ii}(X)}{\alpha_2^2 \lambda_0^2}\right)^{-\frac{1}{2}}. \end{aligned} \tag{17}$$

Rather than using (17), however, the distributional results are used to compute the discretized relative belief ratios as in Example 1. For this $\delta > 0$ is required to determine an appropriate discretization and it will be assumed here that this is the same for all the β_i, although the procedure can be easily modified if this is not the case in practice. Note that such a δ is effectively determined

by the amount that $x_i\beta_i$ will vary from 0 for $x \in R$. Since $x_i \in [-1,1]$ then $|x_i\beta_i| \leq \delta$ provided $|\beta_i| \leq \delta$. When this variation is suitably small as to be immaterial, then such a δ is appropriate for saying β_i is effectively 0. Note that determination of the hyperparameters and δ is dependent on the application.

Again inference can be made concerning $\xi = \Xi(\beta, \sigma^2)$, the proportion of the β_i effectively equal to 0. As in Example 1, however, we can expect bias when the amount of variability in the data is large relative to δ or the prior is too diffuse. To implement the relative belief multiple testing algorithm the quantities (8) and (9) need to be computed to determine q_R and q_A, respectively. The conditional prior distribution of $(b, ||y - Xb||^2)$, given (β, σ^2), is $b \sim N_{k+1}(\beta, \sigma^2(X'X)^{-1})$ statistically independent of $||y - Xb||^2 \sim \text{gamma}((n-k-1)/2, \sigma^{-2}/2)$. So computing (8) and (9) can be carried out by generating (β, σ^2) from the relevant conditional prior, generating $(b, ||y - Xb||^2)$ given (β, σ^2), and using (17).

To illustrate these computations the diabetes data set discussed in Efron, Hastie, Johnstone and Tibshirani (2006) and Park and Casella (2008) is now analyzed. With $\gamma = 0.99$, the values $m_0 = 100, s_1 = 75, s_2 = 200$ were used to determine the prior together with $\tau_0 = 1.05$ determined from the X matrix. This lead to the values $\lambda_0 = 0.48, \alpha_1 = 7.29, \alpha_2 = 13641.35$ being chosen for the hyperparameters. Using the methods developed in Evans and Moshonov (2006), a first check was made on the prior on σ^2 against the data and a tail probability equal to 0.19 was obtained indicating there is no prior-data conflict with this prior. Given no prior-data conflict at the first stage, the prior on β was then checked and the relevant tail probability of 0.00 was obtained indicating a strong degree of conflict. Following the argument in Evans and Jang (2011) the value of λ_0 was increased to choose a prior weakly informative with respect to our initial choice and this lead to choosing the value $\lambda_0 = 5.00$ and then the relevant tail probability equals 0.32.

Using this prior, the relative belief estimates, ratios and strengths are recorded in Table 6. From this it is seen that there is strong evidence against $\beta_i = 0$ for the variables sex, bmi, map and ltg and no evidence against $\beta_i = 0$ for any other variables. There is strong evidence of in favor of $\beta_i = 0$ for age and ldl, moderate evidence in favor of $\beta_i = 0$ for the constant, tc, tch and glu and perhaps only weak evidence in favor of $\beta_i = 0$ for hdl.

As previously discussed it is necessary to consider the issue of bias, namely, compute the prior probability of getting a false positive for different choices of q_R and the prior probability of getting a false negative for different choices of q_A. The value of (8) is 0.0003 when $q_R = 1$ and so there is virtually no bias in favor of false positives and one can feel confident that the predictors identified as having an effect do so. The story is somewhat different, however, when considering the possibility of false negatives via (9). For example, with $q_A = 1$, then (9) equals 0.9996 and when $q_A = 100$ then (9) equals 0.7998. So there is substantial bias in favor of the null hypotheses and undoubtedly this is due to the diffuseness of the prior. The implication is that we cannot be entirely confident concerning those β_i assigned to be equal to 0. Recall, that the first prior proposed lead to prior-data conflict and as such a much more diffuse prior was substituted. The bias in favor of false negatives could be mitigated by

Variable	Estimates	$RB_i(0 \mid X, y)$	Strength
Constant	2	2454.86	0.44
age	−4	153.62	0.95
sex	−224	0.13	0.00
bmi	511	0.00	0.00
map	314	0.00	0.00
tc	162	33.23	0.36
ldl	−20	57.65	0.90
hdl	167	27.53	0.15
tch	114	49.97	0.37
ltg	496	0.00	0.00
glu	77	66.81	0.23

Table 6: Relative belief estimates, relative belief ratios and strengths for assessing no effect for the diabetes data in Example 4..

making the prior less diffuse. It is to be noted, however, that this is an exercise that should be conducted prior to collecting the data as there is a danger that the choice of the prior will be too heavily influenced by the observed data. The real cure for any bias in an application is to collect more data. ■

Next we consider the application to regression with $k+1 > n$.

Example 5. *Non-full rank regression.*

In a number of applications $k+1 > n$ and so X is of rank $l < n$. In this situation, suppose $\{\mathbf{x}_1, \ldots, \mathbf{x}_l\}$ forms a basis for $\mathcal{L}(\mathbf{x}_1, \ldots, \mathbf{x}_k)$, perhaps after relabeling the predictors, and write $X = (\mathbf{1}\ X_1\ X_2)$ where $X_1 = (\mathbf{x}_1\ \ldots\ \mathbf{x}_l)$. For given $r = (X_1\, X_2)\beta_{1:k}$ there will be many solutions $\beta_{1:k}$. A particular solution is given by $\beta_{1:k*} = (X_1(X_1'X_1)^{-1}\ 0)'r$. The set of all solutions is then given by $\beta_{1:k*} + \ker(X_1\, X_2)$ where $\ker(X_1\, X_2) = \{(-B'\ I_{k-l})'\eta : \eta \in R^{k-l}\}, B = (X_1'X_1)^{-1}X_1'X_2$ and the columns of $C = (-B'\ I_{k-l})'$ give a basis for $\ker(X_1\, X_2)$. Given that sparsity is expected for the true $\beta_{1:k}$, it is natural to consider the solution which minimizes $||\beta_{1:k}||^2$ for $\beta_{1:k} \in \beta_{1:k*} + \mathcal{L}(C)$. Using $\beta_{1:k*}$, and applying the Sherman-Morrison-Woodbury formula to $C(C'C)^{-1}C'$, this is given by the Moore-Penrose solution

$$\beta_{1:k}^{MP} = (I_k - C(C'C)^{-1}C')\beta_{1:k*} = (I_l\ B)'\omega_{1:l} \tag{18}$$

where $\omega_{1:l} = (I_l + BB')^{-1}(\beta_{1:l} + B\beta_{l+1:k})$.

From (11) with $\Sigma_0 = \lambda_0^2 I_{k+1}$, the conditional prior distribution of $(\beta_0, \omega_{1:l})$ given σ^2 is $\beta_0 \,|\, \sigma^2 \sim N(0, \sigma^2\lambda_0^2)$ independent of $\omega_{1:l} \,|\, \sigma^2 \sim N_l(0, \sigma^2\lambda_0^2(I_l + BB')^{-1})$ which, using (18), implies $\beta_{1:k}^{MP} \,|\, \sigma^2 \sim N_k(0, \sigma^2\Sigma_0(B))$, conditionally independent of β_0, where

$$\Sigma_0(B) = \lambda_0^2 \begin{pmatrix} (I_l + BB')^{-1} & (I_l + BB')^{-1}B \\ B'(I_l + BB')^{-1} & B'(I_l + BB')^{-1}B \end{pmatrix}.$$

With $1/\sigma^2 \sim \text{gamma}_{\text{rate}}(\alpha_1, \alpha_2)$, this implies that the unconditional prior of the i-th coordinate of $\beta_{1:k}^{MP}$ is $\left(\lambda_0^2\alpha_2\sigma_{ii}^2(B)/\alpha_1\right)^{1/2} t_{2\alpha_1}$.

Putting $X_* = (\mathbf{1}\ X_1 + X_2B')$ gives the full rank model $y \,|\, \beta_0, \omega_{1:l}, \sigma^2 \sim N_n(X_*(\beta_0, \omega'_{1:l})', \sigma^2 I_n)$. As in Example 4 then $(\beta_0, \omega_{1:l}) \,|\, y, \sigma^2 \sim N_l(\omega(X_*, y), \sigma^2\Sigma(X_*)), 1/\sigma^2 \,|\, y \sim \text{gamma}_{\text{rate}}((n + 2\alpha_1)/2, \alpha_2(X_*, y)/2)$ where $\omega(X_*, y) = \Sigma(X_*)X'_*X_*b_*, b_* = (X'_*X_*)^{-1}X'_*y$ and

$$\Sigma^{-1}(X_*) = \begin{pmatrix} n & 0 \\ 0 & (X_1 + X_2B')'(X_1 + X_2B') \end{pmatrix} + \lambda_0^{-2} \begin{pmatrix} 1 & 0 \\ 0 & (I_l + BB') \end{pmatrix},$$
$$\alpha_2(X_*, y) = ||y - X_*b_*||^2 + (X_*b_*)'(I_n - X_*\Sigma(X_*)X'_*)X_*b_* + 2\alpha_2.$$

Now noting that $(X_1 + X_2B')'(X_1 + X_2B') = (I_l + BB')X'_1X_1(I_l + BB')$, this implies $b'_* = (\bar{y}, (I_l + BB')^{-1}b_1)$, where $b_1 = (X'_1X_1)^{-1}X'_1y$ is the least-squares estimate of $\beta_{1:l}$, and

$$\Sigma(X_*) = \begin{pmatrix} n + \lambda_0^{-2} & 0 \\ 0 & (I_l + BB')X'_1X_1(I_l + BB') + \lambda_0^{-2}(I_l + BB') \end{pmatrix}^{-1},$$
$$\omega(X_*, y) = \Sigma(X_*)X'_*X_*b_* = \begin{pmatrix} n\bar{y}/(n + \lambda_0^{-2}) \\ (I_l + BB' + \lambda_0^{-2}(X'_1X_1)^{-1})^{-1}b_1 \end{pmatrix}.$$

Using (18), then $\beta_0 \,|\, y, \sigma^2 \sim N(n(n + \lambda_0^{-2})^{-1}\bar{y}, \sigma^2(n + \lambda_0^{-2})^{-1})$ independent of $\beta^{MP}_{1:k} \,|\, y, \sigma^2 \sim N_k(\beta^{MP}(X, y), \sigma^2\Sigma^{MP}(X))$ where

$$\beta^{MP}(X, y) = \begin{pmatrix} Db_1 \\ B'Db_1 \end{pmatrix}, \quad \Sigma^{MP}(X) = \begin{pmatrix} E & EB \\ B'E & B'EB \end{pmatrix}$$

with $D = (I_l + BB' + \lambda_0^{-2}(X'_1X_1)^{-1})^{-1}$ and $E = ((I_l + BB')(X'_1X_1)(I_l + BB') + \lambda_0^{-2}(I_l + BB'))^{-1}$. The marginal posterior for β_i^{MP} is then given by $\beta_i^{MP}(X, y) + \{\alpha_2(X_*, y)\sigma_{ii}^{MP}(X)/(n+2\alpha_1)\}^{1/2}t_{n+2\alpha_1}$. Relative belief inferences for the coordinates of $\beta^{MP}_{1:k}$ can now be implemented just as in Example 4.

We consider a numerical example where there is considerable sparsity. For this let $X_1 \in R^{n\times l}$ be formed by taking the second through l-th columns of the $(l + 1)$-dimensional Helmert matrix, repeating each row m times and then normalizing. So $n = m(l + 1)$ and the columns of X_1 are orthonormal and orthogonal to $\mathbf{1}$. It is supposed that the first $l_1 \leq l$ of the variables giving rise to the columns of X_1 have $\beta_i \neq 0$ whereas the last $l - l_1$ have $\beta_i = 0$ and that the variables corresponding to the first $l_2 \leq k - l$ columns of $X_2 = X_1B \in R^{n\times(k-l)}$ have $\beta_i \neq 0$ whereas the last $k - l - l_2$ have $\beta_i = 0$. The matrix B is obtained by generating

$$B = \begin{pmatrix} B_1 & 0 \\ 0 & B_2 \end{pmatrix}$$

where $B_1 = (z_1/||z_1|| \cdots z_{l_2}/||z_{l_2}||)$ with $z_1, \ldots, z_{l_2} \overset{i.i.d.}{\sim} N_{l_1}(0, I)$ independent of $B_2 = (z_{l_2+1}/||z_{l_2+1}|| \cdots z_{k-l-l_2}/||z_{l_{k-l-l_2}}||)$ with $z_{l_2+1}, \ldots, z_{k-l-l_2}$ i.i.d. $N_{l-l}(0, I)$. Note that this ensures that the columns of X_2 are all standardized. Furthermore, since it is assumed that the last $l-l_1$ variables of X_1 and the last $k-l-l_2$ variables of X_2 don't have an effect, the form of B is necessarily of the diagonal form given. For, if it was allowed that the last $k - l - l_2$ columns of X_2 were

$k = 10$	Classified Positive	Classified Negative	Total
True Positive	5	0	5
True Negative	1	4	5
Total	6	4	10
$k = 20$	Classified Positive	Classified Negative	Total
True Positive	7	0	7
True Negative	0	13	13
Total	7	13	20
$k = 50$	Classified Positive	Classified Negative	Total
True Positive	7	0	7
True Negative	0	43	43
Total	7	43	50
$k = 100$	Classified Positive	Classified Negative	Total
True Positive	7	0	7
True Negative	0	93	93
Total	7	93	100

Table 7: Confusion matrices for the numerical example in Example 5.

linearly dependent on the the first l_1 columns of X_1, then this would induce a dependence on the corresponding variables and this is not the intention in the simulation. Similarly, if the first l_2 columns of X_2 were dependent on the last $l - l_1$ columns of X_1, then this would imply that the variables associated with these columns of X_1 have an effect and this is not the intention.

The sampling model is then prescribed by setting $l = 10, l_1 = 5, l_2 = 2$, with $\beta_i = 4$ for $i = 1, \ldots, 5, 11, 12$ with the remaining $\beta_i = 0$, $\sigma^2 = 1, m = 2$, so $n = 22$ and we consider various values of $k \geq l$. It is is to be noted that a different data set was generated for each value of k. The prior is specified as in Example 4 where the values $\lambda_0^2 = 4, \alpha_1 = 11, \alpha_2 = 12$ were chosen so that there will be no prior-data conflict arising with the generated data. Also, we considered several values for the discretization parameter δ. A hypothesis was classified as true if the relative belief ratio is greater than 1 and classified as false if it is less than 1. Table 7 gives the confusion matrices with $\delta = 0.1$. The value $\delta = 0.5$ was also considered but there was no change in the results.

One fact stands out immediately, namely, in all of these example only one misclassification was made and this was in the full rank ($k = 10$) case where one hypothesis which was true was classified as a positive. The effect sizes that exist are reasonably large, and so it can't be expected that the same performance will arise with much smaller effect sizes, but it is clear that the approach is robust to the number of hypotheses considered. It should also be noted, however, that the amount of data is relatively small and the success of the procedure will only improve as this increases. This result can, in part, be attributed to the fact that a logically sound measure of evidence is being used. ■

5 Conclusions

An approach to the problem of multiple testing has been developed based on the relative belief ratio. It is argued in Evans (2015) that the relative belief ratio is a valid measure of evidence as it measures change in belief as opposed to belief and, among the many possible candidates for such a measure, it is the simplest with the best properties. One can expect that statistical procedures based on valid measures of evidence will perform better than procedures that aren't as they possess a sounder logical basis. For the multiple testing problem this is reflected in the increased flexibility as well as in the results.

It seems that an appropriate measure of evidence in a statistical problem requires the specification of a prior. While this may be controversial to some, it is to be noted that there are tools for dealing with the subjective nature of some of the ingredients to a statistical analysis such as the sampling model and prior. In particular, there is the process of checking for prior-data conflict after the data is obtained and possibly modifying the prior based upon the idea of weak informativity when such a conflict is encountered. Before the data is actually collected, one can measure to what extent a particular prior will bias the results based upon the particular measure of evidence used. If bias is encountered several mitigating steps can be taken but primarily this will require increasing the amount of data collected. These concepts play a key role in the multiple testing problem.

Acknowledgements

Thanks to Professor Lei Sun for making her notes on multiple testing available.

Bibliography

Baskurt, Z. and Evans, M. (2013) Hypothesis assessment and inequalities for Bayes factors and relative belief ratios. Bayesian Analysis, 8, 3, 569-590.

Cao, Y., Evans, M. and Guttman, I. (2014) Bayesian factor analysis via concentration. Current Trends in Bayesian Methodology with Applications, edited by S. K. Upadhyay, U. Singh, D. K. Dey and A. Loganathan. CRC Press.

Carvalho, C. M., Polson, N. G. and Scott, J. G. (2009) Handling sparsity via the horseshoe. Journal of Machine Learning Research W&CP 5: 73-80.

Efron, B., Hastie, T, Johnstone, I., and Tibshirani, R. (2006) Least angle regression. The Annals of Statistics, 32, 407-499.

Evans, M. (1997) Bayesian inference procedures derived via the concept of relative surprise. Communications in Statistics, 26, 1125-1143.

Evans , M. (2015) Measuring Statistical Evidence Using Relative Belief. Monographs on Statistics and Applied Probability 144, CRC Press.

Evans, M. and Jang, G. H. (2011a) A limit result for the prior predictive. Statistics and Probability Letters, 81, 1034-1038.

Evans, M. and Jang, G. H. (2011b) Weak informativity and the information in one prior relative to another. Statistical Science, 26, 3, 423-439.

Evans, M. and Moshonov, H. (2006) Checking for prior-data conflict. Bayesian Analysis, 1, 4, 893-914.

George, E. I. and McCulloch, R. E. (1993). Variable selection via Gibbs sampling. Journal of the American Statistical Association, 88,881-889.

George, E. I. and McCulloch, R. E. (1997). Approaches for Bayesian variable selection. Statistica Sinica, 7(2), 339–373.

Hastie, T., Tibshirani, R. and Martin Wainwright, M. (2015) Statistical Learning with Sparsity: The Lasso and Generalizations. Monographs on Statistics and Applied Probability 143, CRC Press.

Park, R. and Casella, G. (2008) The Bayesian Lasso. Journal of the American Statistical Association, 103, 681-686.

Rockova, V. and George, E. I. EMVS: The EM approach to Bayesian variable selection. Journal of the American Statistical Association, 109, 506, 828 - 846.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society B., 58, 1, 267-288.

Appendix

Proof of Lemma 4: Let $\Delta(i)$ be the event that exactly i of $A_1, \ldots, A_k \in \mathcal{F}$ occur, so that $\cup_{i=1}^k A_i = \cup_{i=1}^k \Delta(i)$ and note that the $\Delta(i)$ are mutually disjoint. When $l < k$,

$$S_{l,k} = \sum_{\{i_1,\ldots,i_l\}\subset\{1,\ldots,k\}} I_{A_{i_1}\cup\cdots\cup A_{i_l}} = \binom{k}{l}\sum_{i=0}^{l-1} I_{\Delta(k-i)} + \sum_{i=l}^{k-1}\left[\binom{k}{l} - \binom{i}{l}\right] I_{\Delta(k-i)}$$
$$= \binom{k}{l}\sum_{i=0}^{k-1} I_{\Delta(k-i)} - \sum_{i=l}^{k-1}\binom{i}{l} I_{\Delta(k-i)}$$

and $S_{k,k} = I_{A_1\cup\cdots\cup A_k}$. Now consider $\binom{k}{l}^{-1} S_{l,k} - \binom{k}{l-1}^{-1} S_{l,k}$ which equals

$$\frac{1}{\binom{k}{l}} \sum_{\{i_1,\ldots,i_l\}\subset\{1,\ldots,k\}} I_{A_{i_1}\cup\cdots\cup A_{i_l}} - \frac{1}{\binom{k}{l-1}} \sum_{\{i_1,\ldots,i_{l-1}\}\subset\{1,\ldots,k\}} I_{A_{i_1}\cup\cdots\cup A_{i_{l-1}}} \qquad (19)$$

If $l = k$, then (19) equals $I_{A_1\cup\cdots\cup A_k} - \sum_{i=0}^{k-1} I_{\Delta(k-i)} + I_{\Delta(1)} = I_{A_1\cup\cdots\cup A_k} - \sum_{i=0}^{k-2} I_{\Delta(k-i)}$ which is nonnegative. If $l < k$, then (19) equals $\binom{k}{l-1}^{-1} I_{\Delta(k-l+1)} + \sum_{i=l}^{k-1}[\binom{i}{l-1}\binom{k}{l-1}^{-1} - \binom{i}{l}\binom{k}{l}^{-1}] I_{\Delta(k-i)}$ which is nonnegative since an easy calculation gives that each term in the second sum is nonnegative. The expectation of (19) is then nonnegative and this establishes the result.